IEEE 802.11
Wireless Access Method and Physical Layer Specifications

Security Issues for Wireless Networks

March 11, 1991

Robert A. Buaas
The Buaas Corporation
and
Raphael Rom
Sun Microsystems, Inc.

Abstract

The purpose of this paper is to surface the security issues relevant to wireless networking in the IEEE 802 context. Five broad categories are addressed: a) how wireless security is different, b) the privacy requirement, c) resource access control and denial of service, d) authentication, and e) encryption and key management. Unique characteristics of the wireless medium appear to require treatment beyond that provided in other 802 contexts. The discussions stimulated herein could well lead to sufficient resolution of the issues raised.

Introduction

The major issue or question is: should 802.11 deal specifically with security and privacy, or let higher layers do the work? So far, few 802 networks have any machinery for dealing with this subject. Any network security problem was, most likely, resolved in an administrative way by the network manager. The wireless medium makes new demands by virtue of its basic nature. It is these that are the subject of this paper. It is asserted that no wireless product is "complete" without implementing solutions to the security problems the new medium poses.

How wireless differs

The most pervasive assumption made for previous 802 media is that they were private. The user's (owner's) premise was secure from external access and tampering, and so was the network medium. If the user had multiple premises and cared enough, link level encryptors protected the point-to-point links connecting the sites. Further, it was assumed that without physical access and connection to the medium, one could not get the data. Once connected, all the data was fair game. The user's view is, generally, that physical protection provides sufficient privacy. Only administrative controls prevented inappropriate use or tampering with the data (e.g., an offender could be fired and/or prosecuted).

In the wireless environment, these assumptions no longer hold. Radio waves are not limited to a particular set of rooms in an office building. Eavesdropping is likely and could become a popular pastime (e.g., HF/VHF/UHF voice channel scanners). The potential for collecting and exploiting competitor's sensitive business data and intellectual property is enormous considering the minimal effort and expense that may be required, the ECPA (United States Federal Government Electronic Communications Protection Act of 1989) notwithstanding. The medium itself is vulnerable to inadvertent or intentional interference, resulting in reduced levels of service. Malicious jamming could deny access to all or selected users, and/or selectively garble data. Another potential source of interference comes from ether contamination resulting from frequency reuse by emitters radiating power levels higher than those required to maintain the desired communication bit error rate.

Privacy requirements

Privacy deals with data disclosure and integrity among authorized communicating partners. It applies both to the information being transported for the user (transfer data) and the control and management of the communication channel (signaling). The latter contains such information as who talks to whom, when, how often, and at what priority. Every protocol has both aspects. In some situations, only the data being transferred needs protection. In others, compromising the identity of the participants in a conversation leads to considerable damage.

Privacy protection for the transfer data can, in principal, be provided above the MAC layer in the protocol since the data is generated above that level. In the wireless environment, bit error rates are orders of magnitude higher than on previous media, requiring much better error checks. Cryptographic checksums (called message authentication codes) work substantially better than cyclic redundency checks. Faulty packets are discovered and discarded at a very low level, minimizing protocol engine congestion, particularly if the protection mechanism is implemented very efficiently. This same protection might well be applied to the transfer data at no additional cost, eliminating the need to provide it at a higher layer. It is worth noting that most existing applications on the market assume a secure network environment, or at least, do not consider security to be an issue to their successful performance. Imagine the management concern and potential legal liability for the disclosure of routine business data by unsuspecting businesses, and thus the marketability of the offending network products.

Access to resources

Networks are built by having two dedicated resources available, computation for protocol handling and communication channel bandwidth. These are usually scarce resources, requiring every effort at their conservation. Minimizing the waste is usually accomplished at a very low layer in the protocol stack. A jammer might attempt to hog the channel by transmitting a flood of no-operation messages, thereby preventing others the service they legitimately deserve. Denial of service is actually a major issue where physical access to the media cannot be controlled. Access control requires identifying the participants using the channel (who is allowed, and who is to be denied access), which in turn requires a mechanism for authenticating participants.

Another important issue of access control is configuration control. The network architecture should define a "self-organizing" property such that any compliantly-equipped participant could get network service with minimal setup. (For example, each of the attendees at this meeting has a portable computer and desires to participate on the network set up to facilitate the conference.) Put another way, this requires that every network node be complete with the needed hardware and software to accomplish any/all tasks for configuration management when called upon to do so. In self-organizing networks, it is important that those being organized have a guaranteed identity.

Authentication

There are two aspects to this problem. The first is that authentication is essential in any environment where the set of possible physical participants exceeds the set of desired participants (example: two competing businesses within wireless range of one another). The other is that participants in a network must be satisfied that they know that they are in fact talking with the those they intended.

In practice, authentication is necessary in almost any environment where the identity of the user plays a role. Most communication-related applications require authentication (e.g., electronic mail or remote file systems) but very few actually do authenticate their users. Rather, they assume that authentication is being provided by the lower level communication facilities (e.g., that packets are delivered (only) to the intended destination. This is true in wired (static), physically controlled environments, but must be explicitly provided in the wireless environment.

In a wireless network of mobile participants, the self-organizing properties of the network require that the identity of the participant play an important role when service "hand-off" occurs. This might happen as the user moves his network terminal away from the current service point, into range of a "better" service point.

To secure or not to secure

The authors assert that consideration of the above issues leads to a compelling commitment to implement security and privacy mechanisms within MAC/PHY, rather than higher in the protocol stack or not at all. Emphasis should be placed on arriving at an acceptable implementation at the lowest possible cost both in additional hardware and software required, as well as in architectural complexity. The most elegant solution is probably not the most desirable. Having made the case for the need for protection, certain technical issues require consideration.

Technical implementation issues

Authentication requires either a local database within each participant, or an accessible agent whose name and key are globally known. Local databases have the problem of having to be reliably updated whenever an additional participant is added to the network, although certain situations favor such a capability. The agent may be used for storing the database of keys, controlling and accessing the database, or it may be used as a trusted entity for all other communications between participants.

There are two principle ways of implementing secure communication: private-key and public-key cryptosystems. Of these, public-key cryptosystems are more flexible but are computationally more complex. It is thus likely that only the authentication process would use public-keys, while some other private-key block cryptosystem of high computational efficiency would insure the channel privacy. This leads to the further optimization of implementing only the private-key system in hardware, leaving the occasional public-key computations in software. While the use of two cryptosystems is only slightly more complicated, this approach is commonplace since the efficiencies gained are well worth the additional trouble, and no conceptual difficulty is added.

There are three types of communications to be analyzed: direct peer-to-peer (PTP), indirect (Forwarded) PTP, and Broadcast communication.

-Direct PTP requires that both partners have somehow exchanged session keys prior to passing traffic. Data transfer is most efficient but the initial key exchange requires several additional steps.

-Forwarded PTP uses a (trusted) forwarder, capable of decrypting the message from the sender and re-encrypting it for the receiver. Communication is less efficient (in that every message is sent twice on the channel). Key management is simpler because each participant need know only one foreign key--that of the forwarder. If all participants use one forwarder, there might be a channel bandwidth bottleneck there.

-Broadcasts can be sent directly or indirectly (using a forwarder). Key management for direct broadcast could be very costly in that each participant would have to know the broadcast key of every other participant. Use of the forwarder is very efficient, because each participant only need know the broadcast key of the forwarder.

The above schemes must be analyzed in terms of message size, complexity of the protocol, computation involved, channel bandwidth required (number of messages per second, amount of data exchanged), trust, and efficiency. Trust is usually inversely proportional to complexity, suggesting simpler protocols.

Another major issue revolves around the encryption technology to be employed. Which algorithms are appropriate? What cryptographic strength is required? What key length should be used? How often should keys be changed? Can/should any speed or computational efficiency be traded for additional encryption strength? Can the system be utilized internationally without restriction? While careful consideration and selection is needed here, no item is particularly complex.

Key management

There are a number of issues involved in key management. Examples include: a) who generates keys, b) how are they transferred between participants, c) what mechanism does this work, d) who and when are key changes instigated, e) how much trust is required of the various entities, f) what databases are available and where, and g) what and when is synchronization enforced and validated? These issues cannot be formalized until the architecture for authentication is determined. The work of 802.10 may well apply here to a large degree, and should be used in tact wherever possible.

Summary

This paper enumerates a number of issues to be considered in shaping a mechanism for protecting a wireless LAN against security vulnerabilities, data disclosure, corruption, and denial of service. The task of selecting and then clearly specifying the system to be used by 802.11-compliant products appears more arduous than it actually is. The appearance results from the fact that it has not yet been accomplished in the 802 context. There is considerable prior work on which to draw.

The authors believe that an authentication mechanism, including a capability for generating and exchanging working private key, is essential to the functioning of a wireless network. Private-key encryption of some sort is highly recommended to protect the transfer data. Use of the same cryptosystem to verify the quality of the received packets deserves careful consideration.

It is hoped that this contribution stimulates the discussions that will bring the needed mechanisms into being.